IN THE UNITED STATES DISTRICT COURT
WESTERN DISTRICT OF TEXAS
AUSTIN DIVISION

Intellectual Ventures I LLC and Intellectual Ventures II LLC,	
Plaintiffs/Counter-Defendants,	**Civil Action No. 1:19-CV-01075-ADA**
v.	**JURY TRIAL DEMANDED**
VMware, Inc.,	
Defendant/Counter-Plaintiff.	

DEFENDANT VMWARE, INC.'S RESPONSIVE CLAIM CONSTRUCTION BRIEF

TABLE OF CONTENTS

I. INTRODUCTION 1
II. DISPUTED TERMS FROM U.S. PATENT NO. RE44,686 (the "'686 patent") 1
A. "modif[y/ied] [a] resource allocation" / "modify[ing] [the] computer resources allocated to a virtual server" ('686 patent claims 5–7) 1
B. "resource unavailable messages resulting from denied requests to modify a resource allocation" ('686 patent claims 5–7) 3
C. "determination that a virtual server is overloaded" ('686 patent claims 5–7) 4
D. "virtual server" ('686 patent claims 5–7) 6
E. "determining that a second physical host can accommodate the requested modified resource allocation" ('686 patent claims 5–7) 8
F. "component configured to" Means-Plus-Function Terms ('686 patent claim 7) 10
III. DISPUTED TERMS FROM U.S. PATENT NO. RE42,726 (the "'726 patent") 12
A. Terms that overlap with disputed claim terms in the '686 patent 12
B. "resource denials" ('726 patent claims 1, 4–5 and 8) 13
C. "quality of service guarantee" ('726 patent claims 1 and 4) 15
D. Mean-Plus-Function Elements ('726 Patent claims 1, 3, 4, 5, 7) 17
a. "a virtual server resource monitor [communicatively coupled to the first physical host and] configured to monitor resource denials and to send a virtual server overloaded signal in response to the resource denials" ('726 patent claims 1 and 5) // "program code for creating a virtual server resource monitor communicatively coupled to the first physical host and configured to monitor resource denials and, in response to the resource denials, to send a virtual server overloaded signal" ('726 patent claim 4) 17
b. "a virtual server resource modifier [communicatively coupled to the first physical host and] configured to receive the virtual server overloaded signal and, in response to the virtual server overloaded signal, to modify a resource allocation for the virtual server and to send a virtual server resource modification signal" ('726 patent claims 1 & 5); "program code for creating a virtual server resource modifier communicatively coupled to the first physical host and configured to receive the virtual server overloaded signal and, in response to the virtual server overloaded signal, to modify a resource allocation for the virtual server and to send a virtual server resource modification signal" ('726 patent claim 4) 18
c. "a load balanc[ing/er] [module] [communicatively coupled to the plurality of physical hosts and] configured to receive the virtual server resource modification signal and to determine whether the first physical host is overloaded and, in response to a determination that the first physical host is overloaded, to send a physical host transfer signal that indicates a second physical host" ('726 patent claims 1 and 5) // "program code for creating a load balancing module communicatively coupled to the plurality of

physical hosts and configured to receive the virtual server resource modification signal and to determine whether the first physical host is overloaded and, in response to a determination that the first physical host is overloaded, to send a physical host transfer signal that indicates a second physical host" ('726 claim 4) 20

d. "a dynamic virtual server mover [communicatively coupled to the plurality of physical hosts and] configured to receive the physical host transfer signal and, in response to the physical host transfer signal, to transfer the virtual server from the first physical host to the second physical host" ('726 patent claims 1 and 5) // "program code for creating a dynamic virtual server mover communicatively coupled to the plurality of physical hosts and configured to receive the physical host transfer signal and, in response to the physical host transfer signal, to transfer the virtual server from the first physical host to the second physical host" ('726 patent claim 4) 21

e. "the dynamic virtual server mover is further configured to direct the first physical host to store, in the file system, a set of system files for the virtual server and to direct the second physical host to access, from the file system, the set of system files for the virtual server, thereby transferring the virtual server from the first physical host to the second physical host" ('726 claims 3 and 7) 22

IV. DISPUTED TERMS FROM U.S. PATENT NO. 7,949,752 (the "'752 patent") 23

A. "exhausted" ('752 patent claims 1, 9 and 24) 23

B. "consumed" (recited in '752 patent claims 1, 9 and 24) 25

C. "service" ('752 patent claims 1, 3, 9 and 24) 26

D. Means-Plus-Function Terms 28

V. DISPUTED TERMS FROM U.S. PATENT NO. RE43,051 (the "'051 patent") 29

A. "virtual server" ('051 patent claims 1, 3 and 6) 29

B. "physical interface[s]" ('051 patent claims 1 and 3) 32

C. physical interfaces and tunnel identifiers in the storing / receiving / determining / sending terms ('051 patent claims 1 and 3) 34

D. "customer forwarding [table/information]" ('051 patent claims 1 and 3) 36

VI. DISPUTED TERMS FROM U.S. PATENT NO. RE44,818 (the "'818 patent") 37

A. "hierarchical token bucket resource allocation"/ "token" (recited in '818 patent claims 1, 17, 30, 32, 33 and 37–42) 37

B. "enforc[e/ing]", "receiv[e/ing]", "classify[ing]", "compar[e/ing]", "forward[ing]", and "buffer[ing]" ('818 patent claims 1, 17, 30, 32, 33, 37, 38, 39, 42) 40

C. "maintaining a connection over a network fabric" ('818 patent claims 1, 17, 30, 32 and 42) 42

D. "virtual storage network interface layer of an application server" / "virtual network interface layer of an application server"/ "virtual interface layer of an application server" ('818 patent claims 1, 17, 30, 32 and 42) 44

E. "one or more input/output virtualization modules comprising computer-readable instructions operative to cause the one or more processors to" performs functions terms ('818 patent claim 17) 45

VII. CONCLUSION 47

TABLE OF AUTHORITIES

Page(s)

Cases

Bell Atlantic Network Services, Inc. v. Covad Comm'cns Grp., Inc.,
262 F.3d 1258 (Fed. Cir. 2001).....42

In re Berg,
320 F.3d 1310 (Fed. Cir. 2003).....8

Edwards Lifesciences LLC v. Cook Inc.,
582 F.3d. 1322 (Fed. Cir. 2009).....41

GE Lighting Sols., LLC v. AgiLight, Inc.,
750 F.3d 1304 (Fed. Cir. 2014).....24

Glob. Equity Mgmt (SA) Pty. Ltd. v. Expedia, Inc.,
No. 2:16-cv-00095-RWS-RSP, 2016 WL 7416132 (E.D. Tex. Dec. 22, 2016).....46

Honeywell Int'l, Inc. v. ITT Industries, Inc.,
452 F.3d 1312 (Fed. Cir. 2006).....41

Imperium (IP) Holdings, Inc. v. Apple, Inc.,
920 F. Supp. 2d 747 (E.D. Tex. 2013).....10

Inventio AG v. ThyssenKrupp Elevator Am. Corp.,
649 F.3d 1350 (Fed. Cir. 2011).....46, 47

Markman v. Westview Instruments, Inc.,
517 U.S. 370 (1996).....28

Merck & Co. v. Teva Pharm. USA, Inc.,
395 F.3d 1364 (Fed. Cir. 2005).....1, 2

Micro Chem., Inc. v. Great Plains Chem. Co.,
194 F.3d 1250 (Fed. Cir. 1999).....29

Microsoft Corp. v. Multi-Tech Systems, Inc.,
357 F.3d 1347 (Fed. Cir. 2004).....41, 42

Novo Indus., L.P. v. Micro Molds Corp.,
350 F.3d 1348 (Fed. Cir. 2003).....10

O2 Micro Int'l v. Beyond Innovation Techn. Co.,
521 F.3d 1351 (2008).....6, 37, 38

Personal Audio, LLC v. Apple, Inc.,
No. 9:09-cv-111, 2011 WL 11757163 (E.D. Tex. Jan. 31, 2011) 46

Poly-America, L.P. v. API Industries, Inc.,
839 F.3d 1131 (Fed. Cir. 2016).....2

Red Rock Analytics, LLC v. Samsung Elecs. Co.,
No. 2:17-cv-101, 2018 WL 1806859 (E.D. Tex. Apr. 16, 2018) 10, 11

Regents of Univ. of Minnesota v. AGA Med. Corp.,
717 F.3d 929 (Fed. Cir. 2013).....41

RLIS, Inc. v. Allscripts Healthcare Solutions, Inc.,
Nos. 3:12cv208, -209, 2013 WL 3772472 (S.D. Tex. July 16, 2013).....18

Sound View Innovations, LLC v. Facebook, Inc.,
No. 16-cv-116 (RGA), 2017 WL 2221177 (D. Del., 2017).....20

in re TLI Comm'cns LLC Patent Litig.,
87 F.Supp.3d 773 (E.D. Va., 2015) 28

Statutes

35 U.S.C. § 112.....17, 28

35 U.S.C. § 112 ¶ 6..... *passim*

TABLE OF ABBREVIATIONS

Abbreviation	Full Name
PTO	United States Patent and Trademark Office
POSITA	Person of Ordinary Skill in the Art
Snoeren Decl.	Declaration of Alex Snoeren, Ph.D. Regarding Claim Construction
'686 patent	U.S. Patent No. RE 44,686
'726 patent	U.S. Patent No. RE 42,726
'937 patent	U.S. Patent No. 6,985,937
'937 FH	File History of U.S. Patent No. 6,985,937
'752 patent	U.S. Patent No. 7,949,752
'051 patent	U.S. Patent No. RE 43,051
'818 patent	U.S. Patent No. RE 44,818

TABLE OF EXHIBITS
FILED WITH VMWARE'S OPENING CLAIM CONSTRUCTION BRIEF

Exhibit	Title	Dkt. No.
Ex. 1	Parties' Agreed Constructions to Claim Terms	54-1
Ex. 2	Declaration of Alex Snoeren, Ph.D. Regarding Claim Construction	54-2
Ex. 3	'937 Patent File History, Applicant Arguments dated November 17, 2003	54-3
Ex. 4	U.S. Patent No. 5,933,603 to Vahalia et al.	54-4
Ex. 5	'937 Patent File History, Non-Final Office Action dated May 10, 2004	54-5
Ex. 6	USPTO Patent Search of claim term "dynamic virtual server mover"	54-6
Ex. 7	Terms in the '726 patent that overlap with disputed claim terms in the '686 patent	54-7
Ex. 8	The parties' proposals for the terms in the '051 involving multiple recitations of "physical interface(s)" and "tunnel identifier(s)"	54-8
Ex. 9	'752 Patent File History, Response to Office Action dated March 16, 2009	54-9
Ex. 10	'752 Patent File History, Office Action dated May 5, 2009	54-10
Ex. 11	'752 Patent File History, Office Action dated October 30, 2009	54-11
Ex. 12	'752 Patent File History, Response to Office Action dated November 13, 2009	54-12
Ex. 13	'752 Patent File History, Office Action dated March 16, 2010	54-13
Ex. 14	'752 Patent File History, Response to Office Action dated August 20, 2010	54-14
Ex. 15	'752 Patent File History, Response to Office Action dated August 5, 2009	54-15
Ex. 16	'752 Patent File History, Final Office Action dated November 8, 2010	54-16
Ex. 17	'752 Patent File History, Response to Final Office Action dated January 4, 2011	54-17
Ex. 18	Report and Recommendation of United States Magistrate Judge, Intellectual Ventures v. HCC Ins. Holdings, Inc., Case No. 6:15-cv-660 (E.D. Tex. August 26, 2016) ("'752 Patent Magistrate Report")	54-18
Ex. 19	Charles Aulds, Linux Apache Web Server Administration, 39 (2001) ("Aulds")	54-19
Ex. 20	Barry Nusbaum, WebSphere Application Servers: Standard and Advanced Features 45 (1999)	54-20
Ex. 21	Ludmila Cherkasova, FLEX: Design and Management Strategy for Scalable Web Hosting Service, 14–15 (Oct. 1999)	54-21
Ex. 22	October 14, 2003 Amendment & Remarks, U.S. Patent Appl. No. 09/526,980	54-22
Ex. 23	Physical Interface, The IEEE Standard Dictionary of Electrical and Electronics Terms (6th ed., 1996)	54-23
Ex. 24	RE43,051 (U.S. Patent Appl. No. 11/858,091) Patent File History, May 17, 2010 Office Action	54-24
Ex. 25	U.S. Patent No. 6,286,047 ("Ramanathan")	54-25

Exhibit	Title	Dkt. No.
Ex. 26	U.S. Patent No. 6,247,057 ("Barrera")	54-26
Ex. 27	Webster definition of consume	54-27
Ex. 28	'752 Patent List of Disputed Claim Terms	54-28
Ex. 29	BEN LAURIE AND PETER LAURIE, APACHE: THE DEFINITIVE GUIDE 163, 177, 242–43, 295 (1999)	54-29
Ex. 30	*Virtual Server*, MICROSOFT COMPUTER DICTIONARY (5th ed. 2002), p. 555 (IV-VMWARE-00004290)	54-30
Ex. 31	'789 Patent (Orig. Patent) prosecution history, 2009.08.25 Resp. to Office Action at 12	54-31
Ex. 32	Competing Parties' Proposals for the '818 Patent	54-32
Ex. 33	Grotto Networking, available at https://www.grotto-networking.com/BBQoS.html	54-33
Ex. 34	http://tldp.org/HOWTO/Traffic-Control-HOWTO/index.html	54-34
Ex. 35	https://lartc.org/howto/lartc.qdisc.classful.html#AEN1071	54-35
Ex. 36	Linux Advanced Routing & Traffic Control HOWTO, located at http://www.oamk.fi/~jukkao/lartc.pdf	54-36
Ex. 37	Traffic Control HOWTO, Version 1.0.2, Martin A. Brown, located at https://www.tldp.org/HOWTO/html_single/Traffic-Control-HOWTO/	54-37
Ex. 38	Traffic Control using tcng and HTB HOWTO, Version 1.0.1, Martin A. Brown, April 2006, located at http://linux-ip.net/articles/Traffic-Control-tcng-HTB-HOWTO.html	54-38
Ex. 39	Bavier, et al., Operating System Support for Planetary-Scale Network Services, Proceedings of the First Symposium on Networked Systems Design and Implementation (March 2004), located at https://www.usenix.org/legacy/events/nsdi04/tech/full_papers/bavier/bavier.pdf	54-39
Ex. 40	Benita, Kernel Korner - Analysis of the HTB Queuing Discipline Software, available at https://www.linuxjournal.com/article/7562 as of January 26, 2005 and printed in the Linux Journal, March 2005.	54-40
Ex. 41	U.S. Patent 7,161,904 titled: "System and method for hierarchical metering in a virtual router based network switch" to Hussein et al.	54-41
Ex. 42	Bavier et al, Container-based Operating SystemVirtualization: AScalable,High-performance Alternative to Hypervisors, Conference Paper in ACM SIGOPS Operating Systems Review, January 2007, located at http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.1018.1012&rep=rep1&type=pdf	54-42
Ex. 43	Valenzuela, J.L., et al., "A Hierarchical Token Bucket Algorithm to Enhance QoS in IEEE 802.11: Proposal, Implementation and Evaluation, IEEE, vol. 4, Sep. 2004 ("Valenzuela Article")	54-43
Ex. 44	Email from J. Deblois to M. Rueckheim dated March 3, 2020	54-44
Ex. 45	Webster's New Work Telecom Dictionary, Definition of Layer	54-45

TABLE OF EXHIBITS
FILED WITH VMWARE'S RESPONSIVE CLAIM CONSTRUCTION BRIEF

Exhibit	Title
Ex. 46	U.S. Patent No. 6,351,775 ("Yu")
Ex. 47	*Token*, Microsoft Computer Dictionary (5th ed. 2002), p. 532 (IV-VMWARE-00004257)
Ex. 48	U.S. Patent No. 6,976,258 ("Goyal")
Ex. 49	'937 Patent File History, (January 8, 2004 Office Action)

I. INTRODUCTION

For the reasons set forth herein and in VMware's opening brief, VMware requests that the Court adopt VMware's proposed claim constructions for the disputed terms.

II. DISPUTED TERMS FROM U.S. PATENT NO. RE44,686 (THE "'686 PATENT")

A. "modif[y/ied] [a] resource allocation" / "modify[ing] [the] computer resources allocated to a virtual server" ('686 patent claims 5–7)

VMware Proposal	IV Proposal
"modif[y/ied] [a] quality of service guarantee" / "modify[ing] [the] quality of service guarantee of a virtual server" *See also* construction of "quality of service guarantee"	"modif[y/ied] set of functions and features of a physical host(s) used in implementing tasks for each virtual server" / "modify[ing] a set of the functions and features of a physical host(s) used in implementing tasks for each virtual server"

IV's proposal should be rejected because it is contrary to the intrinsic record, including a clear prosecution history disclaimer, and because it is based on a truncated reading of the disputed claim term. Specifically, IV concedes that it is asking this Court to construe the term "**resource allocation**" in accordance with how the specification refers to the term "**resource**." Dkt. No. 53 at 14[1] ("IV's proposal takes its support directly from the intrinsic record which clearly and unambiguously defines 'resource'."). This is improper. A "resource allocation" is not the same thing as a "resource," and IV has provided no explanation to the contrary. *See also* Dkt. No. 54-2 (Snoeren Decl.) at 10 (explaining that a POSITA would have understood these terms to be different). IV's proposal improperly disregards claim language, contrary to Federal Circuit law. *Merck & Co. v. Teva Pharm. USA, Inc.*, 395 F.3d 1364, 1372 (Fed. Cir. 2005) ("A claim construction that gives meaning to all the terms of the claim is preferred over one that does not do so.").

[1] Page citations herein refer to the ECF page numbering unless otherwise noted.

IV's arguments against VMware's proposal are not persuasive. First, IV argues that VMware's proposal "reads out 'resource allocation' entirely and replaces it with 'quality of service guarantee." Dkt. No. 53 at 15. But, in contrast to IV's proposal, which in fact reads out the term "allocation," VMware proposes construing (not reading out) the term based on a clear prosecution history disclaimer by the patentee that equated "resource allocation" with "quality of service guarantee" and "modifying a resource allocation" with "modifying a quality of service guarantee." *See* Dkt. No. 54 at 12–13. This understanding of the claims was expressly acknowledged by the PTO examiner. *Id*. The specification also includes a definitional statement as well as additional clear and consistent statements equating these terms. *Id.* at 13–16. IV's brief fails to address any of this important evidence.

IV also argues against VMware's proposal because the term "quality of service guarantee" is present in the non-limiting preambles of two of the seven claims at issue. Dkt. No. 53 at 15. To the extent that IV is making a claim differentiation argument, that doctrine is inapplicable here. As an initial matter, claim differentiation is merely a presumption, and cannot overcome the clear disclaimer and definitional statements identified by VMware. *Poly-America, L.P. v. API Industries, Inc.*, 839 F.3d 1131, 1137 (Fed. Cir. 2016). Moreover, the procedural history of these reissue patents provides further reason not to apply the presumption. The two independent claims which include a preamble reciting "the computer resources allocated to the virtual server being specified as a quality of service guarantee" ('726 patent claims 1 and 4) are unchanged from the claims as originally issued in the parent '937 patent, which was prosecuted on behalf of the original assignee Ensim. However, the other five independent claims ('726 patent claims 5 and 8; '686 patent claims 5–7) were newly added in the reissue proceeding prosecuted on behalf of Digital Asset Enterprises, an apparent affiliate of IV. Under this procedural posture, it cannot be presumed

that the invention disclosure envisions separate meanings for these terms. Instead, the history suggests that the reissue revisions were an improper attempt to recapture subject matter that had been disclaimed by the original applicant during prosecution.

B. "resource unavailable messages resulting from denied requests to modify a resource allocation" ('686 patent claims 5–7)

VMware Proposal[2]	IV Proposal
"indications that requests by the virtual server for additional resources are either implicitly or explicitly denied, resulting from denied requests to modify a resource allocation" *See also* construction of "modify a resource allocation"	See IV proposals for "resource unavailable messages" and "denied requests to modify a resource allocation" "resource unavailable messages" = "an indication that a request by the virtual server cannot be immediately serviced" "denied requests to modify a resource allocation" = "a request by the virtual server that cannot be immediately serviced"

IV contends that its construction "stays true to the specific sequence of the claim element as a whole." Dkt. No. 53 at 16. As detailed in VMware's opening brief, this is decidedly not the case. Dkt. No. 54 at 16–17.

Instead, IV's proposed construction for "denied requests to modify a resource allocation" is an attempt to re-write the claim as "denied requests ~~to modify a resource allocation~~." IV concedes this point by asking the Court to construe both "resource unavailable messages" and "denied requests to modify a resource allocation" in accordance with "a general definition of 'resource denials'"—a term that is not tied to the claim language of "requests to modify a resource allocation." Dkt. No. 53 at 16–17. IV's proposal should be rejected on at least this basis alone.

[2] As a preliminary issue, IV's opening brief attributes an incorrect proposed construction to VMware for this term. VMware's actual proposed construction is set forth in this table as well as in VMware's opening brief. Dkt. No. 54 at 16–17.

There are other problems with IV's proposal. First, the claim recites "denied requests to **modify a resource allocation**." Both sides have proposed a construction for "modify a resource allocation" (see section II.A). Tellingly, IV's proposal for the present claim language incorporates neither construction; while VMware's proposed construction stays true to the claim language by including the reference to "resulting from denied requests to modify a resource allocation."

Second, IV's proposal for "denied requests to modify a resource allocation" as "a request by the virtual server that cannot be immediately serviced" completely writes out several words in the claim term itself – including "denied," "modify," and "resource allocation." Specifically, IV's proposal does not in any way construe the term "modify." And it does not in any way construe "resource allocation" (or "resource" or "allocation" for that matter). IV's proposal simply *ignores* the presence of these terms in the claim language. Equally problematic, as noted above, IV offers no explanation for why it is proposing near identical constructions for "resource unavailable messages" and "denied requests to modify a resource allocation."

IV's proposal should therefore be rejected and VMware's proposal, which gives due credit to the each and every word in this claim term, should be adopted.

C. "determination that a virtual server is overloaded" ('686 patent claims 5–7)

VMware Proposal	IV Proposal
"determination that an average number of resource denials for a virtual server is beyond a pre-configured threshold" *See also* construction of "resource denials"	Plain and ordinary meaning

IV contends that this term should be given its plain and ordinary meaning because the term and its constituent parts, including the term "overloaded," would have been understood by a POSITA at the time of the invention. Dkt. No. 53 at 23–24. IV's stance on this term is directly contradicted by the fact that IV has proposed a construction different than "plain and ordinary

meaning" for the similar term "virtual server overload signal" claimed in the '726 patent (which includes in common with the present term, the terms "virtual server" and "overload"). Indeed, IV proposed the term "virtual server overload signal" for construction in the first place. IV's stance is also contradicted by the fact that IV proposed the '686 patent claim term – "indication that a first physical host is overloaded" – for construction, and again, proposed a construction different than "plain and ordinary meaning."[3] For the same reasons that IV believed these other terms require construction, the present term merits a construction as well.

Furthermore, as detailed in VMware's opening brief, overload in the context of a physical host is described very differently from overload in the context of a virtual server in the specification. Dkt. No. 54 at 18–19. As such, a POSITA at the time of the invention would have understood that these terms are different. IV argues that VMware's proposal significantly limits the disputed term by importing limitations from an exemplary embodiment. This is not the case. VMware's proposal is drawn from a clear definition of this term in the specification that does not use the word "embodiment" in the sentence (or the paragraph) in which the definition is found:

> A determination is made 220 as to whether a particular virtual server resource is overloaded. The number of times a particular resource denial is received in a time window is averaged using one of a number of well-known techniques. **If the average number of denials is beyond a pre-configured threshold, the virtual server is determined 220 to be overloaded for the corresponding resource**. If the virtual server is not determined to be overloaded, the method continues to monitor 210 virtual server resource denials.

'686 patent, 5:42–50 (emphasis added). IV's opening brief does not consider this portion of the specification and instead cites to unrelated sections of the specification to suggest, incorrectly, that VMware's proposal is importing an illustrative embodiment in the specification into the claims.

[3] The parties have agreed on a construction of the '686 patent claim term "indication that a first physical host is overloaded" as "indication that a first physical host would not support additional resource allocations at that time" and the '726 patent claim term "the first physical host is overloaded" as "the first physical host will not support additional resource allocations at that time."

It bears noting that for the similar term "virtual server overload signal" found in the '726 patent, IV does not propose a plain and ordinary meaning construction, but instead proposes that this term should be construed as "an indication that a virtual server has been or is being denied resources." However, the specification is devoid of any support for this construction, either in a particular embodiment or otherwise. And IV specifically fails to cite any intrinsic support for its proposed construction. By contrast, VMware's proposed construction is not only supported by the specification, but it is consistent across these respective claim terms in the '686 and '726 patents.

D. "virtual server" ('686 patent claims 5–7)

VMware Proposal[4]	IV Proposal
"a process executing on a host computer that accepts communication requests, and that is capable of receiving a quality of service guarantee from a physical host"	plain and ordinary meaning; in the alternative: "a virtual machine that resides on a physical server and uses the physical server's resources but has the appearance of being a separate dedicated machine"

As detailed below in section V.A, VMware's proposal for this term is consistent with a POSITA's understanding of this term (who would have factored in the clear definition from the specification). *See, e.g.,* '686 patent, 3:53–55 ("term 'virtual server' as used herein refers to a virtual server capable of receiving a quality of service guarantee from a physical host"); Dkt. No. 54-2 (Snoeren Decl.) at 13–14, 42–47.

IV's plain and ordinary meaning proposal should be rejected because the parties clearly dispute what the plain and ordinary meaning is. *O2 Micro Int'l v. Beyond Innovation Techn. Co.*, 521 F.3d 1351, 1362 (2008).

[4] As a preliminary issue, IV's opening brief attributes an incorrect proposed construction to VMware for this term. VMware's actual proposed construction is set forth in this table as well as in VMware's opening brief. Dkt. No. 54 at 19.

IV's alternative proposal should be rejected at least because it ignores the specification's clear definition for this term set forth above. It should also be rejected because it inserts the language "virtual machine" into this claim. The term "virtual machine" is not used anywhere in the intrinsic record of these patents, nor is the phrase "appearance of being a separate dedicated machine." Moreover, as discussed further in section V.A, portions of IV's proposed construction are cherry-picked from dictionary support, while other portions lack any support. IV's use of "virtual machine" in its construction is also prejudicial to VMware (whose name stands for "Virtual Machine"-ware). An improper and unsupported construction like this could cause a jury to prematurely judge the merits of IV's infringement allegations despite VMware's defenses.

The intrinsic record also supports VMware's proposal. For example, IV argues that the '937 patent's file history "similarly makes clear that the term 'virtual server' is being used in its customary manner." Dkt. No. 53 at 18. VMware agrees.

Specifically, IV's argument focuses on statements in the '937 patent's file history regarding a prior art reference – U.S. Patent No. 6,351,775 to Yu ("Yu") (Ex. 46). IV argues that the use of the term "virtual servers" by the original applicant in describing Yu amounts to an acknowledgement "that virtual servers are known in the art" and that the term was being used "consistent[ly] with its known meaning, and not some unconventional or heretofore unknown construct [] coining a 'virtual server.'" Dkt. No. 53 at 18. This is an important concession by IV because Yu describes virtual servers in the same manner detailed below in section V.A – i.e. as process-based server applications, such as web servers which can provide for virtual hosting (*i.e.*, servicing requests for multiple network addresses on a single physical host).

For example, Yu's field of invention explains that "[t]he present invention relates generally to providing load balancing across a collection (or cluster) of servers such as proxy servers and

Web servers in the Internet environment." Ex. 46 at 1:21–24 (emphasis added). Yu further explains that "[t]he number of virtual servers is greater than the actual number of servers in the server cluster. The [Domain Name Server] DNS (167) and [Transmission Control Protocol] TCP router can then **dynamically map each virtual server to one of the actual servers in the cluster**." *Id*. at 11:61–64 (emphasis added).

As such, Yu describes virtual servers as a type of **process**, which unlike virtual machines, can be mapped to a physical host. Dkt. No. 54-2 (Snoeren Decl.) at 42–44. The examiner, who is presumed to make informed findings as to the meaning of prior art references to a POSITA,[5] understood Yu the same way, and explicitly equated a virtual server with a **process**:

> - [It is noted that, as taught by Yu at col.11, lines 54-64, each virtual server (which is a process and can be represented by the identifier (e.g., URL) of the object requested) is assigned a class and through the class-to-server mapping, each virtual server is eventually mapped to an actual server residing in a physical host. In other words, there is mapping between a requested object and a physical host (see col.15, lines 21-37)].

Ex. 49 (2004-01-08 Office Action) at 3; *see also id.* at Dkt. No. 54-5 at 5–6 (2004-05-04 Office Action stating same). Indeed, the original patent applicant never disputed this understanding that a virtual server is a type of process.

E. "determining that a second physical host can accommodate the requested modified resource allocation" ('686 patent claims 5–7)

VMware Proposal	IV Proposal
Indefinite	plain and ordinary meaning; in the alternative:

[5] *In re Berg*, 320 F.3d 1310, 1315 (Fed. Cir. 2003) ("As persons of scientific competence in the fields in which they work, examiners […] are responsible for making findings, informed by their scientific knowledge, as to the meaning of prior art references to persons of ordinary skill in the art and the motivation those references would provide to such persons.")

	"determining that a second physical host can accommodate the request(s) by the virtual server that could not be immediately serviced"

IV's opening brief highlights why the antecedent basis problem in this claim language renders the term indefinite. IV argues that "'requested modified resource allocation' ultimately refers back to the denied request to modify a resource allocation." Dkt. No. 53 at 22. The relevant claim language is copied below:

> wherein the determination that a virtual server is overloaded is based on one or more resource unavailable messages resulting from **denied requests to modify a resource allocation**;
>
> [a component configured to] determin[e/ing] that a second physical host can accommodate **the requested modified resource allocation**;

'686 patent, 14:39–45 (emphasis added). As shown, the first part of the claim refers to "denied requests" (i.e., **more than one** denied request), while the second part of the claim refers to a determination of whether a host can accommodate "the requested modified resource allocation" (i.e., **one** request). As such, even accepting IV's explanation of the antecedent basis, ambiguity remains because the claim offers no guidance as to *which one* of the plurality of denied requests to modify a resource allocation that "the requested modified resource allocation" refers back to. Nor does IV cite anything in the specification that would offer guidance to a POSITA on this issue.

Furthermore, IV's alternative proposal that seeks to construe the term "the requested modified resource allocation" the same way as it proposed to construe the language "denied requests to modify a resource allocation" is not helpful. For reasons detailed in section II.B, IV's alternative construction should be rejected because (i) there is no intrinsic support for this interpretation of this claim language, (ii) a "resource unavailable message" / "resource denial" is clearly different from a "requested modified resource allocation", and (iii) IV's proposal for these terms remains at odds with its own proposed constructions of "modified resource allocation" and "modify a resource allocation."

Under Federal Circuit case law, the Court can only correct a mistake in the claim language if "the correction is not subject to reasonable debate based on consideration of the claim language and the specification." *Novo Indus., L.P. v. Micro Molds Corp.*, 350 F.3d 1348, 1354 (Fed. Cir. 2003). IV has not met that standard here. *Imperium (IP) Holdings, Inc. v. Apple, Inc.*, 920 F. Supp. 2d 747, 753 (E.D. Tex. 2013) (finding claims indefinite because there was no antecedent basis for a plural term "pixels" which was argued to refer back to a singular term "pixel"); *see also Red Rock Analytics, LLC v. Samsung Elecs. Co.*, No. 2:17-cv-101, 2018 WL 1806859, at *18 (E.D. Tex. Apr. 16, 2018). As such, because a POSITA would be unable to discern the scope of this language with reasonable certainty, the Court should hold that this claim term is indefinite.[6]

F. "component configured to" Means-Plus-Function Terms ('686 patent claim 7)[7]

IV's arguments are unpersuasive. First, IV argues that VMware cannot show that "the claim "recites function without reciting sufficient structure for performing that function" because the preamble and element preceding the "component" limitations recite structure such as a "virtual server operating in a first physical host" and "one or more processors and one or more memories." Dkt. No. 53 at 26–27.

As a preliminary matter, IV's argument regarding the effect of the preamble on the remainder of the claim is fundamentally contrary to its stance that the preambles of the claims 1 and 4 of the '726 patent (which also recite a "virtual server operating in a first physical host") are not limiting. Dkt. No. 53 at 15. IV cannot have it both ways.

Notwithstanding this issue, IV fails to acknowledge that the claimed "components" are not recited as part of **any** of the claim terms that it contends provide structure to the claims. Instead,

[6] Although VMware initially proposed an alternative construction of this claim term, the parties' respective opening briefs have made it clear that this alternative cannot resolve the indefiniteness issue present in the '686 patent claims.

[7] The parties proposed terms and construction are identified at Dkt. No. 54 at 21–24.

the claim recites broadly, "[a] system" for modifying a virtual server in a physical host, where the components are merely part of the "system." '686 patent, cl. 7. It is of no moment whether "virtual server," "physical host", "processors" or "memories" connote structure when the claims do not require the claimed component terms to be a part of these alleged structures.

IV's reliance on *Fisher-Rosemount Sys., Inc. v. ABB Ltd.* is misplaced. No. 4:18-cv-00178, 2019 WL 6830806 (S.D. Tex. Dec. 12, 2019). *Fisher-Rosemount* did not involve claim language of "a component configured to" perform a function. Instead, *Fisher-Rosemount* found that the claim term "**processor**" connotes structure because, *inter alia*, a processor is a "tangible object that can be purchased and that can perform certain functions even without specific instructions." *Id*. at *16. Additionally, the court noted that claim "describes how the processor interacts with the invention's other components and identifies where the processor is located." *Id*. In contrast, in the present case, the claim language of "a component configured to" clearly doesn't refer to a tangible object that can be purchased and perform functions without specific instructions. Nor does the claim specify *any* relationship between the claimed components and the rest of the system. As such, "a component configured to" is a nonce term—indeed *Fisher-Rosemount* acknowledged that similar terms (such as "element" and "device") "typically do not connote structure." *Id*.

IV's alternative identification of structure disclosed in the specification also fails. First, IV's citation to Figure 1, 2:63–3:4, 3:59–4:8, 5:7–28, 5:42–62[8] and multiple structures "Dynamic Resource Configuration Module 100; Physical Hosts 160A-C; Virtual Servers 162A-G" for each claimed "component," is unsupported attorney argument. A POSITA reading these sections would not understand them to disclose support for the claimed functions. Dkt. No. 54-2 (Snoeren Decl.)

[8] In the parties' various exchanges, and despite requests from VMware, IV has never cited to portions of the specification that it contends perform the claimed function. IV disclosed these citations for the first time in its opening brief.

at 14–20. In particular, the claimed functions here, for the three claimed component terms, generally relate to (i) receiving an indication that a first host is overloaded based on a specific determination of whether a virtual server is overloaded based on one or more resource unavailable messages resulting from denied requests to modify a resource allocation; (ii) determining that the second physical host can accommodate the requested modified resource allocation; and (iii) generating a host transfer signal indicating a second physical host. '686 patent, cl. 7. IV's specification citations completely fail to disclose a software algorithm for performing the first two claimed functions. For example, the specification fails to discuss *anything* related to the claim language of determining that a virtual server is overloaded based on one or more resource unavailable messages resulting from denied requests to modify a resource allocation.[9] As such, this claim is indefinite.[10]

III. DISPUTED TERMS FROM U.S. PATENT NO. RE42,726 (THE "'726 PATENT")

A. Terms that overlap with disputed claim terms in the '686 patent

Similar / Overlapping Terms		
Term	**VMware Proposal**	**IV Proposal**
"modify[ing] a resource allocation for the virtual server" / "modifying [the] computer resources allocated to a virtual server" ('726 patent claims 1, 4-5 and 8)	"modify[ing] a quality of service guarantee for the virtual server" / "modifying [the] quality of service guarantee of a virtual server" *See also* construction of "quality of service guarantee"	"modify[ing] a set of the functions and features of a physical host(s) used in implementing tasks for each virtual server" / "modify[ing] a set of the functions and features of a physical host(s) used in implementing tasks for each virtual server"

[9] Additionally, IV's citations fail to discuss anything related to the claim language of determining that a virtual server is overloaded based on one or more resource unavailable messages resulting from denied requests to modify a resource allocation.

[10] As to the third claimed function, IV's specification citations miss the mark because they fail to include citations to the algorithm in the specification describing "Dynamic Virtual Server Mover 140." While column and lines 5:7–28 (cited by IV) discuss the virtual server mover at a high level, the algorithm corresponding to the Figure 6 (and described in the specification at 12:1–28) is necessary to provide structure to this means-plus-function term.

"virtual server overloaded signal" ('726 patent claims 1, 4-5 and 8)	"signal indicating that an average number of resource denials for a virtual server is beyond a pre-configured threshold" *See also* construction of "resource denials"	"an indication that a virtual server has been or is being denied resources"
"virtual server" ('726 patent claims 1-11)	"a process executing on a host computer that accepts communication requests, and that is capable of receiving a quality of service guarantee from a physical host"	plain and ordinary meaning; in the alternative: "a virtual machine that resides on a physical server and uses the physical server's resources but has the appearance of being a separate dedicated machine"

For the reasons set forth in sections II.A, II.C, and II.D that respond to IV's arguments regarding these terms, the Court should adopt VMware's proposed constructions of these terms.

B. "resource denials" ('726 patent claims 1, 4–5 and 8)

VMware Proposal	IV Proposal
"indications that requests by the virtual server for additional resources are either implicitly or explicitly denied"	"an indication that a request by the virtual server cannot be immediately serviced"

IV contends that its proposed construction "is a direct quotation from the portion of the specification where the invention as a whole […] is described" while "VMware's proposed construction is merely one narrow example of a single preferred embodiment." Dkt. No. 53 at 15. Not so.

First, as detailed in VMware's opening brief, the parties have reached substantial agreement regarding the first part of the construction of "resource denials" – i.e., "indications that requests by the virtual server." Dkt. No. 54 at 24–25. But, a subtle, but important, distinction for the present term is that VMware incorporates plural nouns in the agreed language as "indications…" to account for the fact that the claim term resource **denials** is a plural noun. In

contrast, IV improperly seeks to alter this term to singular tense as "an indication." Dkt. No. 54 at 24–26.

Second, unlike VMware's proposed construction, IV's proposal omits the word "resource" from the claim term, improperly attempting to broaden out the requests by the virtual server for additional resources to simply any request by the virtual server. These attempts to substantially re-write the claims are contrary to black letter claim construction law and should be rejected.

Third, VMware's proposal is taken from a definitive statement in the specification of "resource denials," which explains that resource denials "**are**" indications that requests by the virtual server "**for additional resources [are] either implicitly or explicitly denied**." '726 patent, 7:51–54 (emphasis added). By contrast, IV's proposal has been selectively taken from language that discusses what a resource denial "**may refer to**." '726 patent, 2:55–61 (emphasis added). Moreover, IV relies upon a section of the specification that refers to a specific embodiment. *See* '726 patent at 2:52–61 ("In one embodiment … [a] resource denial may refer to any request by the virtual server that cannot be immediately serviced").

Fourth, IV's argument that VMware's proposed construction is based on an illustrative embodiment is not correct. Neither the sentence nor the paragraph in the specification from which VMware's proposed construction is drawn uses the language "embodiment." *See* '726 patent, 7:50–64. And the sentence from which VMware's proposal is drawn makes it clear that the characterization of resource denials was intended to be definitional. By contrast, as noted above, the phrase "[i]n one embodiment" is used in the same paragraph, and indeed, in the sentence preceding the one from which IV draws its proposed construction. Under IV's argument, this fact undermines its own proposed construction. For these reasons, IV's proposed construction should be rejected and VMware's proposed construction should be adopted.

C. "quality of service guarantee" ('726 patent claims 1 and 4)

VMware Proposal	IV Proposal
"information that specifies a guaranteed amount of an assigned resource, and that can be dynamically increased/modified"	"a guaranteed resource allotment which can be dynamically increased/modified"

IV takes issue with two portions of VMware's proposed construction – "information that specifies" and "an assigned resource" – and contends that "the intrinsic record is devoid of any such limitation." Dkt. No. 53 at 20.

As to the language "information that specifies," IV contends that the claim term "quality of service guarantee" is not information and is different from "quality of service guarantee information." This argument obfuscates what is a simple issue. A quality of service guarantee is unquestionably information. The specification is replete with examples that make it clear that this must be the case.

For example, the summary of the invention states that "[t]he present invention **dynamically adjusts the quality of service guarantees** for virtual servers based upon the resource demands experienced by the virtual servers." '726 patent, 2:38–40 (emphasis added). It is unclear, and IV has not explained, how it would be possible to dynamically adjust quality of service guarantees if they did not refer to information. The specification also states that "[t]he term "virtual server" as used herein refers to a virtual server capable of **receiving a quality of service guarantee** from a physical host." '726 patent, 3:43–45 (emphasis added). It is unclear how a virtual server would receive a quality of service guarantee from a physical host if it was not information. The specification goes on to state that "[a] **resource allocation for a virtual server is specified as a "quality of service guarantee"** for that particular server. Each **physical host stores quality of service guarantees** for the virtual servers it hosts." '726 patent, 4:39–42 (emphasis added). Again, it is not clear how a resource allocation could be specified as a quality of service guarantee or how

a physical host could store quality of service guarantees if these terms did not refer to information. These instances are provided by way of example and without limitation, as there are other such examples in the specification. As such, IV's arguments as to the language "information that specifies" should be rejected and VMware's proposed construction, which properly incorporates this language into the construction, should be adopted.

As to the language "assigned resource," IV argues that "nowhere in the specification is a quality of service guarantee described as being limited to 'an assigned' resource." Dkt. No. 53 at 21. This is not true. The specification states:

> A customer's virtual server is typically **assigned a fixed level of resources**, **corresponding to either a fixed percentage of the capacity of a particular physical host** (for example, the operating system may be instructed to allocate twenty percent of the central processing unit cycles to process A and two percent to process B) **or a fixed number of units** (for example, the operating system may be instructed to allocate X cycles per second to process A and Y cycles per second to process B).

'726 patent, 2:5–13 (emphasis added). This directly parallels later language that states:

> In one embodiment, each individual virtual server 162 has a different quality of service guarantee. Different **quality of service guarantees are implemented by allocating different amounts of the resources of each physical host machine** 160 to servicing each of the virtual servers 162. **Physical host 160 resources may be allocated as percentages of the resources of a particular physical host 160, or as a particular number of units within a physical host 160** (for example, the operating system may be instructed to allocate X cycles per second to process A and Y cycles per second to process B)

'726 patent, 3:66–4:8 (emphasis added). As such VMware's proposed language finds clear support in the specification. By contrast, the language "guaranteed resource allotment," and sub-phrases "resource allotment" and "allotment" (and for that matter, other versions of allotment such as "allot" or "allotted"), in IV's proposed construction are not found *anywhere* in the specification. As such, for the reasons IV's itself sets forth, its proposed construction should be rejected in favor of VMware's proposed construction.

D. Mean-Plus-Function Elements ('726 Patent claims 1, 3, 4, 5, 7)

VMware contends that each of the terms identified in Dkt. No. 54 at 30–33 should be construed under § 112 as means-plus-function terms. IV in contrast argues that terms recited in the preamble such as "network system," "physical hosts," "virtual server," and "computer resources," "give context to the environment in which the 'virtual server resource monitor' is claimed." Dkt. No. 53 at 28. As discussed in section II.F, this is directly contradicted by IV's argument that the preambles of the '726 patent claims are not limiting. Furthermore, as discussed in section II.F, none of these terms are recited as part of **any** of the claim terms that IV contends provide structure. Instead, the '726 patent claims 1, 3, 4, 5, 7 recite a "network system," "computer program," or "system" comprising these separate claim terms and do not provide context or structure **any** of these terms. VMware's opening brief provides additional discussion regarding why these are black-box terms that fail to carry structural meaning to a POSITA. *See e.g.,* Dkt. No. 54 at 27–33; Dkt. No. 54-2 (Snoeren Decl.) at 24–42. Additional reasons why the Court should not adopt IV's proposals for these terms are discussed below.

a. "a virtual server resource monitor [communicatively coupled to the first physical host and] configured to monitor resource denials and to send a virtual server overloaded signal in response to the resource denials" ('726 patent claims 1 and 5) // "program code for creating a virtual server resource monitor communicatively coupled to the first physical host and configured to monitor resource denials and, in response to the resource denials, to send a virtual server overloaded signal" ('726 patent claim 4)

Every single instance of this term in the specification describes what the black-box "monitor" *does*, not what it *is*. Because the term itself does not carry structural meaning to a POSITA (either based on or independent of the intrinsic record), and because the claims functionally recite what this black-box module is "configured to" do, this term invokes § 112 ¶ 6. Also, as described in VMware's opening brief, the term "program code" does not on its own

connote structural meaning to a POSITA, and is not discussed anywhere in the specification.[11] As such, and because the claims functionally recite what this black-box term is "configured to" do, this term invokes § 112 ¶ 6.

IV's alternative identification of structure disclosed in the specification also fails. IV cites to Figure 1, Figure 3, 4:64–5:4, 5:44–65[12] and multiple structures including "dynamic resource configuration module 100 includes, among other things, virtual service resource modifier, which in combination with physical hosts 160A-C monitors resource denials and sends a virtual server overload signal" as providing structure for the "virtual server resource monitor." However, IV specifically does not cite to the "virtual server resource monitor," even though it separately contends that this term is sufficiently described in the specification to connote structure for a POSITA. Dkt. No. 53 at 29–30. This argument does not make sense.

By contrast, VMware points to the software algorithm in the specification that describes the functions of the "Virtual Server Resource Monitor 110," which is the same as the claim term, as the disclosure of structure. VMware also points to a more complete discussion of this term in the specification than IV, which specifically omits column and lines 7:41–9:46 despite its disclosure of a software algorithm that performs the claimed function. IV never explains how or why VMware's identification of structure unreasonably limits the claimed functionality. If anything, IV's limited citations that fail to include the relevant algorithm, are incomplete.

b. "a virtual server resource modifier [communicatively coupled to the first physical host and] configured to receive the virtual server overloaded signal

[11] IV's reliance on the *RLIS*, *Eolas*, and *Aloft* opinions does not warrant a different result. Dkt. No. 53 at 29. Each of these cases were decided under the overruled "strong presumption" standard. *Compare RLIS, Inc. v. Allscripts Healthcare Solutions, Inc.*, Nos. 3:12cv208, -209, 2013 WL 3772472, at *14 (S.D. Tex. July 16, 2013) ("[T]he presumption flowing from the absence of the term 'means' is a strong one that is not readily overcome") with *Williamson*, 792 F.3d at 1349 ("expressly overrule the characterization of that presumption as 'strong'").

[12] Again, IV disclosed these citations for the first time in its opening brief.

and, in response to the virtual server overloaded signal, to modify a resource allocation for the virtual server and to send a virtual server resource modification signal" ('726 patent claims 1 & 5); "program code for creating a virtual server resource modifier communicatively coupled to the first physical host and configured to receive the virtual server overloaded signal and, in response to the virtual server overloaded signal, to modify a resource allocation for the virtual server and to send a virtual server resource modification signal" ('726 patent claim 4)

For the reasons discussed above in section III.D.a, the Court should construe the claim language of "a virtual server resource modifier … configured to … [perform a function]" and "program code for [performing a function]" as invoking 35 U.S.C. § 112 ¶ 6. Every single instance of this term in the specification describes what the black-box "modifier" *does*, not what it *is*.

IV's alternative identification of structure disclosed in the specification also fails. IV cites to Figure 1, Figure 2A, Figure 4, 4:64–5:20[13] and multiple structures including "dynamic resource configuration module 100 includes, among other things, virtual service resource modifier, and in combination with physical hosts 160A-C and virtual servers 162A-G receives virtual server overload signals and signals a resource modification is needed" as providing structure for the "virtual server resource modifier." This argument and the identification of other multiple other components than "Virtual Server Resource Modifier 120" for performing the functions of the claimed "virtual server resource modifier" does not make sense.

By contrast, VMware points to the software algorithm in the specification that describes the functions of the "Virtual Server Resource Modifier 120," which is the same as the claim term, as the disclosure of structure. VMware also points to a more complete discussion of this term in the specification than IV, which specifically omits column and lines 9:47–10:52 despite its disclosure of a software algorithm that performs the claimed function. IV never explains how or

[13] Again, IV disclosed these citations for the first time in its opening brief.

why VMware's identification of structure unreasonably limits the claimed functionality. IV's limited citations that fail to include the relevant algorithm, are incomplete.

c. "a load balanc[ing/er] [module] [communicatively coupled to the plurality of physical hosts and] configured to receive the virtual server resource modification signal and to determine whether the first physical host is overloaded and, in response to a determination that the first physical host is overloaded, to send a physical host transfer signal that indicates a second physical host" ('726 patent claims 1 and 5) // "program code for creating a load balancing module communicatively coupled to the plurality of physical hosts and configured to receive the virtual server resource modification signal and to determine whether the first physical host is overloaded and, in response to a determination that the first physical host is overloaded, to send a physical host transfer signal that indicates a second physical host" ('726 claim 4)

IV contends that the terms "load balance[r] module" and "program code for creating" the same are "specific and well-known construct[s] with a structure known by those of skill in the art." Dkt. No. 53 at 32. IV provides no support for this argument. A POSITA would not understand these terms to refer to well-known structures. Dkt. No. 54-2 (Snoeren Decl.) at 33–34. Indeed, IV even admits that these terms relate broadly to something "typically implemented in software." Dkt. No. 53 at 32. IV's descriptions of this module show that it invokes § 112(6). *Id*. at 32–33. *See, e.g., Sound View Innovations, LLC v. Facebook, Inc*., No. 16-cv-116 (RGA), 2017 WL 2221177, at *2 (D. Del., 2017) ("For software patents claiming a function that a general purpose computer cannot perform, the specification must disclose an algorithm."). Indeed, similar to the previous terms, every single instance of this term in the specification describes what the black-box "module" *does*, not what it *is*. Because the term itself does not carry structural meaning to a POSITA (either based on or independent of the intrinsic record), and because the claims functionally recite what this black-box module is "configured to" do, this term invokes § 112 ¶ 6.

IV's alternative identification of structure disclosed in the specification also fails. IV cites to Figure 1, Figure 5, Figure 6, 4:64–5:20, 10:53–11:20[14] and multiple structures including "dynamic resource configuration module 100 includes, among other things, load balancer, and in combination with physical hosts 160A-C and virtual servers 162A-G receives virtual server resource modification signal and determines whether the physical host is overloaded and sends a host transfer signal if it is" as providing structure for the "load balancing module." This argument and the identification of other multiple other components than "Physical Host Load Balancing Module 130" for performing the functions of the claimed "load balancing module" does not make sense.

By contrast, VMware points to the software algorithm in the specification that describes the functions of the "Physical Host Load Balancing Module 130," which is effectively the same as the claim term, as the disclosure of structure. VMware also points to a more complete discussion of this term in the specification than IV, which specifically omits, among other cites, column and lines 11:21–52 despite its disclosure of a software algorithm that performs the claimed function. IV never explains how or why VMware's identification of structure unreasonably limits the claimed functionality. IV's limited citations that fail to include the relevant algorithm, are incomplete.

d. "a dynamic virtual server mover [communicatively coupled to the plurality of physical hosts and] configured to receive the physical host transfer signal and, in response to the physical host transfer signal, to transfer the virtual server from the first physical host to the second physical host" ('726 patent claims 1 and 5) // "program code for creating a dynamic virtual server mover communicatively coupled to the plurality of physical hosts and configured to receive the physical host transfer signal and, in response to the physical host transfer signal, to transfer the virtual server from the first physical host to the second physical host" ('726 patent claim 4)

[14] Again, IV disclosed these citations for the first time in its opening brief.

For the reasons discussed above in section III.D.a, the Court should construe the claim language of "a dynamic virtual server mover … configured to … [perform a function]" and "program code for [performing a function]" as invoking 35 U.S.C. § 112 ¶ 6. Every single instance of this term in the specification describes what the black-box "mover" *does*, not what it *is*.

IV's alternative identification of structure disclosed in the specification also fails. IV cites to Figure 1, Figure 6, 4:64–5:20, 6:20–31, 6:39–59, 12:1–10[15] and multiple structures including "dynamic resource configuration module 100 includes, among other things, dynamic virtual service mover and in combination with physical hosts 160A-C" as providing structure for the "dynamic virtual server mover." This argument and the identification of other multiple other components than "Dynamic Virtual Server Mover 140" for performing the functions of the claimed "dynamic virtual server mover" does not make sense.

By contrast, VMware points to the software algorithm in the specification that describes the functions of the "Dynamic Virtual Server Mover 140," which is the same as the claim term, as the disclosure of structure. VMware also points to a more complete discussion of this term in the specification than IV, which specifically omits portions of column and lines 11:63–12:23 despite its disclosure of a software algorithm that performs the claimed function. IV never explains how or why VMware's identification of structure unreasonably limits the claimed functionality. IV's limited specification citations that fail to include the relevant algorithm, are incomplete.

e. "the dynamic virtual server mover is further configured to direct the first physical host to store, in the file system, a set of system files for the virtual server and to direct the second physical host to access, from the file system, the set of system files for the virtual server, thereby transferring the virtual server from the first physical host to the second physical host" ('726 claims 3 and 7)

[15] Again, IV disclosed these citations for the first time in its opening brief.

For the reasons discussed above in section III.D.d, and in VMware's opening brief, the Court should construe the language which recites "dynamic virtual server mover … configured to … [perform a function]" as invoking 35 U.S.C. § 112 ¶ 6.

IV. DISPUTED TERMS FROM U.S. PATENT NO. 7,949,752 (THE "'752 PATENT")

A. "exhausted" ('752 patent claims 1, 9 and 24)

VMware Proposal	IV Proposal
"unavailable for reuse"	"used up to the allotted or pre-determined amount"

IV makes two primary arguments for this term. Neither are persuasive. And tellingly, IV has failed to address the clear prosecution disclaimer supporting VMware's proposed construction. *See, e.g.,* Dkt. No. 53 at 9–10.[16]

First, IV argues that the intrinsic record is "replete with evidence supporting IV's proposed construction." Dkt. No. 53 at 9 (citing various passages from the '752 patent's specification). Not so. The '752 patent only uses the word "exhausted" once—specifically in the context of halting consumption of a service resource when "the amount held by an agent is exhausted." '752 patent at 25:30–34. None of the other sections cited by IV reference the concept of "exhausted." Moreover, there are no passages in the specification that overcome IV's clear and unambiguous prosecution disclaimer labeling the "exhausted" claim term as the **exact opposite** of **reusable** resources. Dkt. No. 54-14 (2010-08-20 Response) at 14 (emphasis added).

[16] Furthermore, IV's argument that the Court should follow Magistrate Judge Mitchell's prior report and recommendation ("R&R") for this term to promote "uniform treatment of claim construction" is inapposite. Dkt. No. 53 at 9. That R&R was merely a "recommendation" and was not adopted by the District Court Judge before the case settled. Indeed, **both parties** objected to the Magistrate's recommendations before settling. Because the claims were not ultimately construed by the District Judge, there exists no risk of inconsistent rulings on the terms.

The bulk of IV's remaining argument is that IV's proposal reflects the overall goal of the invention. Again, this is not the case. VMware's proposal is consistent with the invention disclosure, while IV's proposal incorrectly interprets a particular embodiment and then seeks to read this interpretation into the claims. *See, e.g.*, *GE Lighting Sols., LLC v. AgiLight, Inc.*, 750 F.3d 1304, 1309 (Fed. Cir. 2014).

The only paragraph that discusses the concept of exhaustion ('752 patent, 25:20–34) illustrates why VMware's proposal is correct. This paragraph is written with reference to "method 800" which, in relevant part, includes steps 810 and 820. *Id.* at 25:20; *see also id.* at Fig. 17 (depicting method 800). Step 810 performs an evaluation of an assigned "service permission" in order to "identif[y] the amount of a service resource 25 which is allocated to agent 22." *Id.* at 24:56–58; *see also id.* at 18:19–21 ("a service permission can specify a pre-authorized amount of service resource 25 which is allowably consumed by the respective agent"). After executing the service, step 820 then "decrement[s] the amount allocated to agent 22 by the amount actually used." *Id.* at 25:16–19; *see also* Fig. 17 ("decrement A_h by A_u"). VMware's proposal is consistent with this disclosure because it recognizes that once the service's resources (A_h) have been fully decremented/used up, they are not reusable without issuance of an additional service permission assigning a new A_h amount of service resources.

IV's proposal, in contrast, is inconsistent with this description. IV's proposal contemplates that the service resource should be available again after the customer "has released the resource." Dkt. No. 53 at 5. But method 800 does not include a step where the amount of service resource A_h is **incremented** by an amount of service resource released. Nor is that process discussed anywhere else in the '752 patent. Nor is IV correct in arguing that under VMware's proposal "[o]ne customer using an amount of a service and service resource would permanently make that

service and service resource unavailable to subsequent customers." *Id*. The claims don't recite that the entirety of a service, and/or a service resource, applicable to multiple customers is exhausted. Instead, they recite a different customer-specific concept, of "a service resource **configured to be consumed by the network-based agent**" where "**an amount** of the service resource is exhausted." *See, e.g.,* '752 patent cl. 7. As discussed above, and in VMware's opening brief, the specification also repeatedly and consistently describes this process as customer specific.

IV's proposal is also inconsistent with the two examples of service resources discussed with connection to method 800: "an amount of storage space" (for use with a voicemail service provider) and an amount of "connect time" (for a long distance service provider). '752 patent at 25:20–34. For example, claim 1 becomes incomprehensible when using IV's proposals for the "exhausted" and "consumption" terms, reading as: "wherein an amount of the service resource is ~~exhausted~~ used up to the allotted or pre-determined amount upon being ~~consumed~~ used by the network-based agent." As written, the user's apportioned resource would be "used up" immediately upon being used. This does not make sense, and it is directly contrary to IV's own articulation of the overall goal of the invention: allowing a customer to pay for computing power, time and access on an as needed basis. Dkt. No. 53 at 9–10. VMware's proposals, in contrast, render the claim viable, reading as: "wherein an amount of the service resource is ~~exhausted~~ unavailable for reuse upon being ~~consumed~~ used up by the network-based agent."

B. "consumed" (recited in '752 patent claims 1, 9 and 24)

VMware Proposal	IV Proposal
"used up"	"used"

As articulated in its opening brief, VMware's proposal for this term is supported by narrowing statements made during prosecution and a definitional statement provided in the patent specification. As evidenced by these statements, "consumed" cannot simply mean "use" as IV

proposes—indeed, the claim language was amended to employ both "using" and "consumed" evidencing that the two words have different meaning. *See, e.g.*, Dkt. No. 54-12 (2009-11-13 Amendment) at 5–6. The applicant then used this narrowing amendment to argue around prior art that disclosed an agent configured to use a resource. *Id.*, at 11 ("Humpleman does not teach or suggest 'a service and a service resource **configured to be consumed by the agent** . . . wherein the service resource is **exhausted** after it is consumed by the agent," as recited by claims [sic] 86, 92, 94, 106, and 107.") (emphasis retained).

Indeed, IV fails to substantively address the narrowing and definitional statements in the intrinsic record (also identified in VMware's opening brief). Instead, IV only points to three areas of the '752 patent, none of which provide a definitional statement of the term "consumed" or otherwise disclose anything other than "consumed" and "used" being included in sentences using the terms as alternative options. *See,* Dkt. No. 53 at 10.

IV cannot escape the fact that the terms "exhausted" and "consumed" are intertwined. VMware's proposed constructions are consistent with the specification and file history and also are consistent with one another. On the other hand, IV's constructions for "exhausted" and "consumed" when read together make the claim language incomprehensible.

C. "service" ('752 patent claims 1, 3, 9 and 24)

VMware Proposal	IV Proposal
"An application that is used by an agent on behalf of a principal"	"Network functionality available to agent(s)/network-based agent(s)"

VMware's proposal is consistent with the specification's disclosure that an "agent uses the service on behalf of a principal" and that the services "comprise one or more software applications providing various capabilities that are available to a principal." '752 Patent at 3:3–5; 10:17–19;

see also id. at 10:19–21 ("Each service 24 may be utilized by one or more agents 22 in order to perform their respective tasks.").

IV's proposal, in contrast, does not help clarify the claim scope. IV cites to the abstract's disclosure that "an agent is operable to utilize a service within the network system"—but IV appears to only cite this section as support that an agent used a service in a **network system**. This clarification, however, is not necessary because the claims already recite that the agent uses a service in a **network system**. *See, e.g.*, '752 patent claim 1 ("means, including the network-based agent, for using a service"). Also, IV's proposed language of "network functionality" could be interpreted as relating to functionality of a network, such as signal processing and transmission—a concept that is inconsistent with the '752 patent's more general description of service applications that utilize a network system but are not network functionality itself, such as "an e-mail service, a voice mail service, a paging/facsimile service, an address book and calendar service, and a business news and stocks information service." '752 patent at 13:25-28.

Nor is IV correct in arguing that the patent discloses examples of the claimed service as "sub-systems" and not "applications" (Dkt. No. 53 at 11). Instead, the specification actually describes "sub-systems" (e.g., "mass storage subsystem[s] of tapes or disk drive" *see* '752 patent at 13:49) as something that ***supports*** a service, but not as the service itself. *See id.* at 13:20–26 ("At least a portion of the sub-systems in computer-based system 30 may support one or more services 24").

Finally, IV's own dictionary definition contradicts its proposal. This definition states that "services" are "specialized, **software-based** functionality provided by network servers." Dkt. No. 53 at 12 (emphasis added). While IV emphasized that the definition includes the word

"functionality," the preceding "specialized, software-based" portion of the definition (i.e., an application) is glossed over.

D. Means-Plus-Function Terms

VMware's proposals for these "means" terms identify "the corresponding structure … described in the specification" as required by 35 U.S.C. § 112. For example, for the "means for mediating an interaction between the means for using the service and the service" term in claim 3, VMware proposes that the structure is a "service wrapper (26) as described in [the '752 patent at] 16:22–38." The specification citation refers to a passage titled "Service Wrapper (Details))" which provides a description of the structure for "service wrapper (26)." This section describes that "service wrapper (26)" is composed of "a converter 48 and a monitor 50." *Id*. at 16:28–29.

In contrast, IV's proposal ("service wrapper (26)"), should be rejected because it does not identify the corresponding structure described in the specification. IV apparently intends to leave it up to the jury to determine whether the disclosed structure for this means term should be interpreted to require, e.g., the "converter 48" and "monitor 50a" as disclosed in the specification. This is improper—construing claim scope is the province of judges. *Markman v. Westview Instruments, Inc*., 517 U.S. 370, 372 (1996). Moreover, imposition of IV's vague proposal would render the claim indefinite. *Compare in re TLI Comm'cns LLC Patent Litig.*, 87 F.Supp.3d 773, 800 (E.D. Va., 2015) ("simply disclosing a black box that performs the recited function is not a sufficient explanation of the algorithm required to render the means-plus-function term definite") (quotations omitted).

VMware did not "cherry-pick[]" portions of the specification ignoring other structure for this term as IV argues. Dkt. No. 53 at 7. Most of IV's additional citations merely discuss a "service wrapper (26)" in functional terms devoid of structure. *See* '752 patent at 3:20–27 (function: "cooperat[ing] with the agent server to mediate interaction"); 17:43–46 (function: "control[ing]

access"); 18:49–54 (no statement of structure or function). The only possible exception is IV's citation to the '752 patent at 25:1–24. Dkt. No. 53 at 13. This section describes an algorithm for service wrapper 26 including steps 802-820. To the extent this Court finds this disclosure as structure, VMware offers a revised proposed construction of "service wrapper (26) as described in 16:22–38 and 25:1–24." *See Micro Chem., Inc. v. Great Plains Chem. Co.*, 194 F.3d 1250, 1258 (Fed. Cir. 1999) ("When multiple embodiments in the specification correspond to the claimed function, proper application of § 112, ¶ 6 generally reads the claim element to embrace each of those embodiment.").

Finally, the parties dispute the structure for the term "means for monitoring an amount of the service resource used by the network-based agent" of claim 4. VMware's proposal is that the "means for monitoring" is a "monitor as described in '752 patent at 16:50–61" and should be adopted. IV, in contrast, argues that the structure should be taken up to the software package of the "service wrapper 26" which, as described above, is described as including a monitor and a converter. However, IV fails to explain how the converter is at all involved in the claimed function and should be rejected for this reason alone. The converter does not perform the function of monitoring, instead, it is used to "convert between a computer language (or instruction set) used within agent server 20 and a computer language (or instruction set) used within the respective service 24." *See, id*. at 16:30–33.

V. DISPUTED TERMS FROM U.S. PATENT NO. RE43,051 (THE "'051 PATENT")

A. "virtual server" ('051 patent claims 1, 3 and 6)

VMware Proposal	IV Proposal
"a process executing on a host computer that accepts communications requests"	Plain and ordinary meaning, or alternatively: "virtual machine(s) that reside(s) on a physical server and use(s) the physical server's resources but has/have the appearance of being a separate dedicated machine(s)"

IV is right that the term "virtual server" is used in the '051 patent (and in the '686 and '726 patents) consistent with its customary usage at the time of the invention. But IV is wrong that the customary use at the time would have meant a virtual machine (which IV defines as a functional simulation of a computer and its associated devices) instead of the application servers actually described in the intrinsic record. There is no intrinsic support for IV's contention that virtual server means simulated machines, like computers. Instead, consistent with VMware's position, the intrinsic record describes the virtual server of the invention as process based.

For instance, the '051 patent explains that instead of dedicating individual physical host computers to individual customers, a "service provider may utilize one physical host computer to provide commercial host services to multiple customers" using "***a server application*** executing on a single physical host [that] can be programmed to process requests made to multiple network addresses." '051 patent at 2:47–53 (emphasis added). The specification describes this to mean "each customer is assigned a network address (or domain name), and is provided with resources on a single, physical host computer, effectively sharing the host with other customers." *Id.*

The '051 patent further explains how to create these private virtual servers by pointing to an incorporated reference that issued as U.S. Patent No. 6,976,258. *Id.* at 4:64–67. The '258 patent in turn describes a server as "a process, executing on a dedicated physical services client [that] services client requests for a single network address (physical host) only," Ex. 48 at 3:23–25, and explains that these processes can provide for virtual hosting (*i.e.*, servicing requests for multiple network addresses on a single physical host) by "creat[ing] child processes to service the requests," *id.* at 1:24–44. This ability for a single application to create child processes to handle client requests on a plurality of network addresses is what provides for the *virtual* aspect of the

virtual servers of the invention, so that the application can function as multiple servers. *See id.* at 1:45–2:5; Dkt. No. 54-2 (Snoeren Decl.), 43–44.

In contrast to the many discussions of process-based virtual servers in the intrinsic record, there is no support whatsoever for IV's incorrect assertion that virtual server means virtualized computer hardware.[17] IV mistakenly points to disclosures about the virtual servers using resources on the host machine as evidence that these virtual servers are somehow partitioning resources for use in virtualized hardware. The specification, however, is clear that the physical host's resources are being allocated to server *applications*, not virtualized hardware. It explains, for example, that "virtual hosting" involves "a server application executing on a single physical host," '051 patent at 2:47–53, such that each customer "is provided with resources on a single, physical host computer, effectively sharing the host with other customers," *id.* at 2:54–64. In other words, IV's citations to discussions about sharing physical resources on a single host apply to process-based virtual hosting, and do not indicate that the virtual servers of the invention are virtualized hardware, as IV suggests.

IV's citation to two cherry picked dictionary definitions does not trump the intrinsic record's support of VMware's proposal. Moreover, a more thorough review of the extrinsic record reinforces the conclusion that the invention describes process-based virtual servers, not simulated hardware. For instance, Ben and Peter Laurie's book, *Apache: The Definitive Guide* (1999), describes how to configure Apache virtual servers and discusses virtual servers in terms of processes. Dkt. 54-29 at 7, 10, 13–14, 16; Dkt. No. 19 at 6 (emphasis added) ("Apache is an example of a preforking server," which "means that the main server ***starts a pool of processes*** to

[17] *See, e.g.,* Dkt. No. 53 at 47 (IV arguing: "a virtual server, i.e., a software abstraction of physical hardware").

handle client requests, rather than forking a new process for each incoming request."). Moreover, the process-based Apache virtual server is discussed in terms of preforking—which is exactly how the intrinsic record for IV's patent discusses virtual servers of the invention. Ex. 48 ('258 patent) at 1:45–2:5; *see also* Dkt. No. 54-2 (Snoeren Decl.) at 42–45.

Given the background provided in the specification, including about the function of forking child processes to service client requests on multiple network addresses, a process-based server application like Apache is what a POSITA would have understood the term virtual server to mean at the time of the invention. As a result, the Court should reject IV's arguments that attempt to read virtualized hardware, like virtual machines, into the claims, and instead give the term virtual server its plain meaning as evidenced by the intrinsic and extrinsic record: a process executing on a host computer that accepts communications requests.

B. "physical interface[s]" ('051 patent claims 1 and 3)

VMware Proposal	IV Proposal
"hardware that provides a point of communication between two or more devices"	Plain and ordinary meaning.

IV argues that physical interface should be given its plain and ordinary meaning, but IV avoids the actual issue: the parties have a fundamental dispute about what the plain and ordinary meaning of physical interface is. IV has steadfastly refused to confirm whether it agrees (or disagrees) that the term "physical" in "physical interface" refers to hardware. IV's opening brief continues to avoid taking a position on whether physical means hardware, and instead focuses on whether the described interface provides a point of communication between two or more devices.

This argument is a red herring because it ignores the crux of the parties' dispute. Moreover, IV is wrong that a physical interface is not a point of communication between two or more devices. Independent claims 1 and 3 both describe the method as "providing private network services . . . in

a location remote from private network users." Reading claims 1 and 3, there is no doubt that the physical interface is providing a point of communication between at least the remote location and the private network users. IV incorrectly tries to sidestep this by pointing out that the tunnel switch may have only one physical interface. But even if true, that does not change that the claims require the physical interface receive transmissions from, and send transmissions to, other devices. *See, e.g.*, '051 patent at claim 1 ("receiving . . . a transmission on a physical interface" and "sending the transmission . . . on the determined physical interface"). There is no indication, even in IV's proposed embodiment where the tunnel switch has only one physical interface, that the physical interface would receive transmissions from itself and send transmissions to itself, particularly given the explanation in the preamble that the transmissions relate to providing network services between network users and a remote location. In other words, IV's argument on this claim term is at best a distraction from the actual dispute between the parties.

Again, the real dispute between the parties is whether the physical interface is hardware, and IV is wrong that the Court should simply disregard this dispute under the guide of applying plain and ordinary meaning. The Court should rather give real meaning to the patentee's choice to use "physical interface" instead of just "interface". Considering how the word physical is used in the specification, how a POSITA would understand that term in the claims is clear: it differentiates the interface as hardware, not a virtual or logical device. For example, the specification uses physical to differentiate between the "physical host" and the "virtual servers" running on it. *See, e.g.*, *id.* at 2:47–64, 4:35–38, 4:56–65, 7:1–11, 7:47–48, 8:4–5. It also explains that the tunnel is a logical construct by differentiating it from a "physical connection." *See id.* at 10:17–20. Although IV sidesteps this entirely, given this context, a POSITA would have understood that the term physical interface refers to hardware; and as explained above, more specifically hardware that

provides a point of communication between two or more devices. The Court should therefore adopt VMware's proposal and reject IV's call for the Court to avoid resolving this fundamental dispute between the parties.

C. physical interfaces and tunnel identifiers in the storing / receiving / determining / sending terms ('051 patent claims 1 and 3)

VMware's proposal adds clarity to the relevant claims, which otherwise risk confusing the jury through recitals of multiple physical interfaces and tunnel identifiers at various steps. *See, e.g.*, '051 patent at cl. 1. For example, some physical interfaces and tunnel identifiers are stored in a customer lookup table. *Id.* Other physical interfaces and tunnel identifiers are stored in customer forwarding tables. *Id.* A transmission is received on a physical interface containing a tunnel identifier, and later, a transmission is sent on a physical interface with a tunnel identifier. *Id.* The claims as written are vague as to which tunnel identifiers and which physical interfaces are being used at each step. But the specification provides clarification: the physical interfaces and tunnel identifiers associated with the receiving step (and thus the customer lookup tables) are the *incoming* physical interfaces and the *incoming* tunnel identifiers, whereas the physical interfaces and tunnel identifiers associate with the sending step (and thus the customer forwarding tables) are the *outgoing* physical interfaces and the *outgoing* tunnel identifiers. *See id.* at 11:30–32, 11:64–12:1, 12:10–20, 12:30–33, 12:59–63, 13:5–11 Figs. 8–9.

IV is wrong to argue that VMware "attempt[s] to limit the disputed terms to a single directionality requirement." Dkt. No. 53 at 51. VMware's proposal does not speak to whether a tunnel is uni- or bi-directional, and instead merely clarify which of the "physical interfaces" and "tunnel identifiers" are being recited. *See* Dkt. No. 54-8. These proposals are necessary because the claims vaguely recite more than one "***a*** physical interface" and tunnel identifier. *See, e.g.*, '051 patent at cl. 1.

IV is also wrong to argue that VMware's proposal contradicts the specification's teaching of a tunnel switch with "one or more physical interfaces." *See* Dkt. No. 53 at 51 (emphasis in original). There is no implication in VMware's proposal that the *incoming* physical interface and the *outgoing* physical interface cannot be the same. This would be intuitive to a POSITA, who would understand that typical network cards are bidirectional and can function both as an outgoing interface and as an incoming interface. IV is also wrong to criticize VMware's proposal for using the word "pair." *See* Dkt. No. 53 at 51. VMware's proposals do not use the word "pair". *See* Dkt. No. 54-8.

VMware's proposal clarifies these terms in a manner consistent with the claim language and specification. The *incoming* tunnel identifiers and the *incoming* physical interfaces are associated with the *receiving* step in the claims. There is no reasonable reading of the claims where the receiving step is not dealing with an incoming transmission. Likewise, the *outgoing* tunnel identifiers and *outgoing* physical interfaces are related to the *sending* step in the claims. There is no reasonable reading of the claims where the sending step is not dealing with an outgoing transmission. The Court should reject IV's alternative suggestion of using "first" and "second" because those terms further obfuscate these elements in the claims. In addition, IV only provides this alternative as to the "tunnel identifier" terms, which would not resolve the ambiguities caused by the "physical interface" terms. Finally, "first" and "second" are not used in the specification to describe these elements. As explained above, the specification consistently uses incoming and outgoing for this purpose.

Ultimately, there is ambiguity in the claims and IV has not identified any problem inherent to using the labels in the specification to clarify which components are being used at which steps. To the contrary, it appears IV prefers to maintain the ambiguity in the claim language. If so, the

Court should hold the claims to be indefinite. Alternatively, the Court should adopt VMware's proposal and use the terms "incoming" and "outgoing" consistently with their use in the specification to clarify the ambiguities in the claim language regarding the physical interfaces and tunnel identifiers used in the various steps of the claim.

D. "customer forwarding [table/information]" ('051 patent claims 1 and 3)

VMware Proposal	IV Proposal
Plain and ordinary meaning	"table(s) containing [a set/sets] of customer specific forwarding information" / "set(s) of customer specific forwarding information"

VMware proposes that no construction of these terms are necessary. The language is not complex and would not be confusing to a jury.

The Court should not adopt IV's proposal. First, IV's proposal for "customer forwarding table" defines the contents of this table broadly (as sets of customer specific forwarding information) without reference to information associating network addresses with outgoing physical interfaces and outgoing tunnel identifiers that is required by other language in the claims. *See, e.g.,* '051 patent at cl. 1 ("the customer forwarding tables associating network addresses with physical interfaces and tunnel identifiers . . .); *see* section IV.C herein (discussing how the concept of "outgoing" interfaces and identifiers relate to the customer forwarding table term).

Second, there is no reason to add IV's proposed language of "specific" or "set/sets" to these claim terms. IV argues that this language is necessary to clarify that there is there is a "single, 'correct' customer forwarding table/information accessible to each customer." Dkt. No. 53 at 44–45. But this argument ignores that there is already language in the claims reciting determination of "the correct" table/information. *See e.g.,* '051 patent at cl. 1 ("determining ***the correct*** customer forwarding table from . . ."); *id.* at cl. 3 ("determining ***the correct*** customer forwarding information from . . .").

Third, it is appropriate for the Court to issue a plain and ordinary meaning construction for this term. The parties apparently do not dispute that the claim language, in the context of the entire claim, require tables (or information) that contain customer specific forwarding information segregated by customer. *See O2 Micro*, 521 F.3d at 1362 (not required to construe limitations if there is no fundamental dispute).

VI. DISPUTED TERMS FROM U.S. PATENT NO. RE44,818 (THE "'818 PATENT")

A. "hierarchical token bucket resource allocation"/ "token" (recited in '818 patent claims 1, 17, 30, 32, 33 and 37–42)

VMware Proposal	IV Proposal
the specific class-based scheduling algorithm known in the art as the "hierarchical token bucket"	Plain and ordinary meaning

IV argues that "there is no need to construe" these terms ("HTB") beyond their plain and ordinary meaning. Dkt. No. 53 at 37. However, resolution of this dispute is necessary because the parties dispute what the plain and ordinary meaning is.

IV's brief (and infringement contentions) confirm that it intends to argue that the plain and ordinary meaning of HTB does not actually mean the HTB algorithm that was well-known to POSITAs, but instead is a "non-limiting" term that could encompass other QoS or scheduling mechanisms. For instance, IV argues that the "claimed invention involves 'a two-tier hierarchical QoS management process . . . employed in a virtual I/O server' in which 'network fabric resources are allocated in a hierarchical arrangement,'" that this "hierarchical resource allocation" can be implemented "using scheduling and queuing methods such as hierarchal token bucket." Dkt. No. 53 at 36. IV's brief repeatedly suggests that the term "HTB" is used in the claims in a "non-limiting" manner. *Id*. at 36–37.[18]

[18] Yet, IV fails to identify the "intended meaning of the term," which is the reason the parties are at an impasse as to the construction of this term. Because IV failed to provide a construction that

But, the claims **are** expressly limited to use of the HTB, and do not extend to any "two-tier hierarchical QoS management process" as IV seems to suggest. This is confirmed in both the claims and the specification. The claims themselves uniformly use the term "hierarchical token bucket" – and not any other QoS or scheduling mechanisms. The specification confirms that the inventors knew that HTB was a **specific one** of many QoS or scheduling mechanisms, and that the inventors intended the term HTB to have its well-known meaning – the specific class-based scheduling algorithm known in the art as the "hierarchical token bucket."[19] The inventors were aware of other QoS and scheduling mechanisms, and chose to limit their claims to just one of them – the HTB. Limiting the claims to HTB in this manner was critical to getting them allowed. Dkt. No. 54-31 at 13 (arguing that HTB was different than other "hierarchical or tree structures for storing resource reservations"). IV should be held to this understanding of the claims.

IV's remaining criticisms of VMware's proposal are not persuasive. IV argues that the articles VMware cites show that "multiple different implementations of a hierarchical token bucket algorithm were known" and that therefore the term "should not be limited to a specific, extrinsic definition as VMware proposes." Not so. Each of these articles show that HTB has one and only one meaning: the specific class-based scheduling algorithm known in the art as the "hierarchical token bucket." It is a proper noun, just like the Statue of Liberty. As Dr. Snoeren explains, and

---

resolves the parties' dispute, the Court should deny enter VMware's proposed construction. *See O2 Micro*, 521 F.3d at 1361.

[19] The specification repeatedly and consistently identifies HTB as a specific QoS mechanism at least 12 times. *See, e.g.* '818 patent at 9:61–65; 9:66–10:2 ("QoS manager 414 … maintain[s] a scheduling mechanism, **such as a HTB** scheduling mechanism, that controls whether packets are forwarded for further processing or enqueued on fabric receive buffer"); 10:15–16 ("Hierarchical token bucket can be considered as a class based scheduling mechanism."); 11:49–50 "using a QoS mechanism **such as hierarchical token bucket** (HTB)"); 17:7–10 ("The SAN QoS manager 424, using mechanism **such as HTB**, determines if the application server 102 associated for the write data has sufficient tokens to transmit the write data to the SANI/O systems.").

explained above, this is how the patent itself uses the term. The extrinsic evidence merely confirms that the HTB was well known in the art, and that the patent itself was referring to that common usage of the term and not anything else.[20] Nor does the article referenced by IV as its Exhibit Q support IV's position. This article refers to HTB as a proper noun consistent with VMware's proposal. This article does not discuss different implementations of HTB, and instead merely discusses how to use HTB in various use cases such as Linux and WLAN.

Regarding the term "token," IV does not criticize VMware's construction or offer its own. Its brief, however, characterizes it as a "standalone" term and suggests IV will argue the term has a meaning that applies more broadly than tokens used in "hierarchical token bucket." *See* Dkt. No. 53 at 37–38. But the term token is not a standalone term – it is used specifically in connection with the HTB. The claims confirm this. For instance, they include steps of "classifying" the packets "relative to the hierarchical token bucket resource allocation to determine a current amount of tokens available," "comparing" the size of the packets "to the current amount of tokens available," and "forwarding" the packets "if the current amount of tokens available are sufficient." *See, e.g.* Claim 42. This confirms that VMware's proposal is correct.

IV has identified extrinsic evidence regarding the term token that suggests it intends to rely on some other meaning for the term. Ex. 47 (Microsoft Computer Dictionary, Fifth Ed., (2002) p. 522) ("A unique structured data object or message that circulates continuously among the nodes of a token ring and describes the current state of the network"); *id*. ("Any nonreducible textual

[20] IV also argues that VMware's ignores the significance of ". . . resource allocation [of bandwidth]" which is a mischaracterization of facts. Dr. Snoeren explains that hierarchical token bucket and token "refer to the very specific method of allocating bandwidth resources referred to in the art as a hierarchical token bucket, or "HTB." *See* Dkt. No. 54-2 (Snoeren Decl.) at 48. Moreover, the claims themselves use "hierarchical token bucket resource allocation of bandwidth" as the antecedent to "hierarchical token bucket resource allocation." *See, e.g.*, '818 patent at 17:51–52, 17:61–62.

element in data that is being parsed, for example, the use in a program or a variable name, a reserved word, or an operator."). This would be contrary to the discussion of token in the claims and specification, and should therefore be rejected.

B. "enforc[e/ing]", "receiv[e/ing]", "classify[ing]", "compar[e/ing]", "forward[ing]", and "buffer[ing]" ('818 patent claims 1, 17, 30, 32, 33, 37, 38, 39, 42)

VMware Proposal	IV Proposal
"enforcing . . . across the physical [storage network] interface of the virtual I/O server" "receiv[e/ing] in the virtual I/O server" "classify[ing] in the virtual I/O server" "compar[e/ing] in the virtual I/O server" "forward[ing] in the virtual I/O server" "buffer[ing] in the virtual I/O server"	Plain and ordinary meaning (for each)

These limitations constitute the HTB steps of the claim (hereinafter, the "HTB Steps"). Beginning with the summary of the invention, the patent makes clear that "the present invention" requires that these limitations be performed **in the virtual I/O server**. '818 patent at 1:66–2:3 ("**the present invention** provides methods and apparatuses directed to managing quality of service (QoS) **in virtual input/output (I/O) servers**." *Id*. at 8:19–21 ("**The present invention** manages QoS of I/O subsystems **in virtual I/O servers** …."). The Federal Circuit has repeatedly held that statements regarding "the present invention" limit the claims and confirm that, in the '818 patent, the claimed steps are required to be performed by the virtual I/O server and not anywhere else. *Honeywell Int'l, Inc. v. ITT Industries, Inc.*, 452 F.3d 1312 (Fed. Cir. 2006) ("We agree with the district court that the claim term 'fuel injection system component' is limited to a fuel filter. . . On at least four occasions, the written description refers to the fuel filter as 'this invention' or 'the

present invention'"); *Edwards Lifesciences LLC v. Cook Inc.*, 582 F.3d. 1322, 1329 (Fed. Cir. 2009) ("Here, the specification frequently describes an 'intraluminal graft' as "the present invention" or "this invention," indicating an intent to limit the invention to intraluminal devices."); *Regents of Univ. of Minnesota v. AGA Med. Corp.*, 717 F.3d 929, 936 (Fed. Cir. 2013).

And in fact, the claims are written from the perspective of the virtual I/O server. The independent claims all include the HTB Steps (enforcing, receiving, classifying, comparing, forwarding, and buffering). The summary of the invention and the entire specification "repeatedly and consistently" confirm that these functions are performed only by the virtual I/O server. Dkt. No. 54 at 54–55. The Federal Circuit has consistently held that where, as here, the patentee has repeatedly and consistently explained its invention, the claims should be so limited. *Microsoft Corp. v. Multi-Tech Systems, Inc.*, 357 F.3d 1347 (Fed. Cir. 2004).[21]

IV argues that VMware's proposals improperly "reads out" preferred embodiments and seeks to "redefine structural relationships" among other claim elements. IV argues that VMware's constructions contradict the specification's supposed disclosure that the "virtual I/O server" is "not necessarily physically distinct from the application server(s) and I/O subsystems that it connects." Dkt. No. 53 at 39. This argument misses the point.[22] As explained above, the "present invention" requires that the HTB Steps (e.g. classifying, buffering, etc.) be performed **by the virtual I/O server** and not the claimed application server or anything else. IV points to nothing in the specification where these functions are performed by anything other than the virtual I/O server.

[21] The Federal Circuit held that "claims must be interpreted in light of the specification" and "repeatedly and consistently" describing a feature of the invention is not "limited to describing a preferred embodiment, but more broadly describes the overall invention." *Id.* at 1347–48.
[22] It is also incorrect. IV misleadingly quotes the "Background" section of the patent. The patent does **not** disclose a virtual I/O server that is physically distinct from the application server.

Not only are VMware's constructions consistent entire specification, they are required by the inventors' clear statements as to what they invented.

C. "maintaining a connection over a network fabric" ('818 patent claims 1, 17, 30, 32 and 42)

VMware Proposal	IV Proposal
"maintaining a connection between the physical interface of the application server and the physical interface of the virtual I/O server over a network fabric"	Plain and ordinary meaning

IV argues that VMware is "attempting to insert additional structural limitations "physical interface of the application server" and "physical interface of the virtual I/O server." Dkt. No. 53 at 40. But, the specification specifically discloses that it is the physical interface of the application server that connects to the network fabric in Figure 2. Dkt. No. 54 at 56. And, the specification repeatedly and consistently refers to the connection with the network fabric as a physical connection. *Id*. ("I/O fabric PHY interface 202 generally refers to the hardware interface or interconnection to the I/O switch fabric . . . Virtual I/O server 106 connects to the I/O switch fabric 104 through I/O fabric interface 110 such as Infiniband ports."). The Federal Circuit has held that "when a patentee uses a claim term throughout the entire patent specification, in a manner consistent with only a single meaning, he has defined that term by implication." *Bell Atlantic Network Services, Inc. v. Covad Comm'cns Grp., Inc.*, 262 F.3d 1258, 1271 (Fed. Cir. 2001). This Court should reject IV's invitation to overturn well-established Federal Circuit precedent and adopt VMware's proposed construction.

IV argues that the "virtual I/O server need not be physically distinct from the application servers and associated virtual network interface(s)" and supports its argument with the following quote from the specification:

> Similarly, **virtual network interface**, in one implementation, emulates an Ethernet NIC. In one implementation, this driver plugs in at the bottom of the

> network stack and provides an Internet Protocol address bridged by the Virtual I/O server 106 onto a LAN.

Dkt. No. 53 at 41; '818 patent at 4:9–13 (emphasis added). But this evidence supports VMware's position. As shown in annotated Figure 2 below, the "virtual network interface" does not connect with the "I/O switch Fabric"; the physical interface of the application server (PHY 202) connects to the I/O switch fabric:

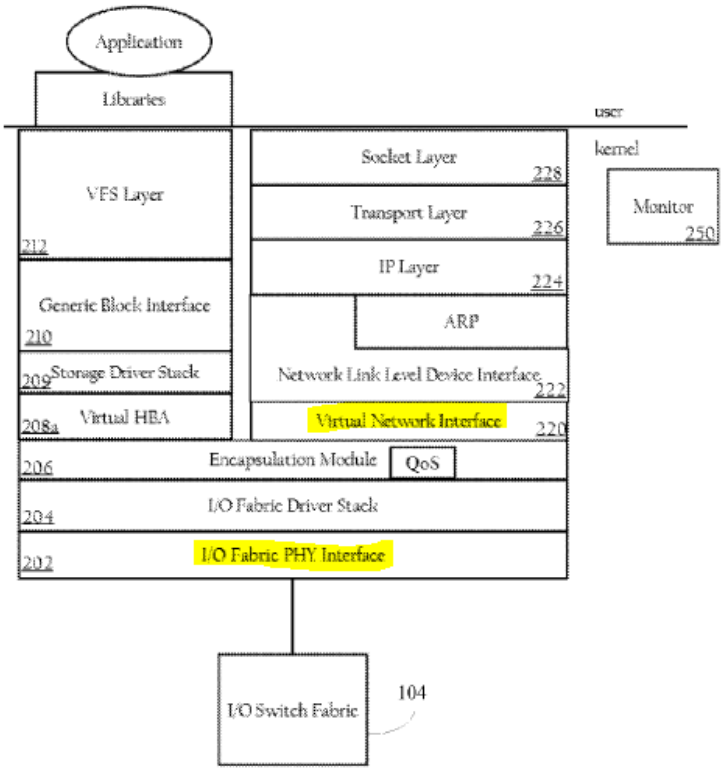

Fig. 2

There is no support in the '818 patent for IV's position that the scope of the claim term can have a broader meaning under the guise of a "plain and ordinary meaning" construction. Especially when the inventor repeatedly and consistently described the invention as maintaining a connection between the physical interface of the application server and the physical interface of the virtual I/O server over a network fabric. The Court should adopt VMware's proposed construction.

D. "virtual storage network interface layer of an application server" / "virtual network interface layer of an application server"/ "virtual interface layer of an application server" ('818 patent claims 1, 17, 30, 32 and 42)

VMware Proposal	IV Proposal
"a virtual storage network interface to higher layers of the virtual node in an application server" / "virtual network interface layer to higher layers of the virtual node in an application server" / "virtual interface layer to higher layers of the virtual node in an application server"	Plain and ordinary meaning

IV identified what it means by "plain and ordinary meaning" for the first time in its opening brief. That is, IV argues that "virtual [network/storage network] interface layer of an application server" means "interface layers (e.g., virtual network interface 220, virtual HBA 208a) that emulate layers of a networking or storage protocol stack" of an application server. Dkt. No. 53 at 35. Given this proposal, VMware proposes the following meaning which is consistent with the intrinsic evidence and IV's proposed construction: "virtual interface layer (e.g., virtual network interface 220, virtual HBA 208a) of an application server that emulates a physical device or a virtual device relative to a physical device." The only dispute between the proposed constructions above is: what does the virtual interface layer emulate? The specification provides the answer:

> The **virtual hostbus adapter (HBA), emulating a physical HBA**, receives SCSI commands for a given device and passes them to the virtual I/O server 106 over the I/O switch fabric. Similarly, **virtual network interface, in one implementation, emulates an Ethernet NIC**.

'818 patent at 4:6–10.

> **HBA module 208a emulates a physical hostbus adapter** relative to the native operating system executed on the application server 102.

Id. at 7:41–43.

> In other implementations, **the virtual HBA module 208 can be configured to emulate a virtual block device relative to the generic block interface**.

Id. at 7:65–67. The specification repeatedly and consistently disclose that the virtual interface layer emulates a physical device (e.g., physical HBA, Ethernet NIC) or a virtual device relative to

a physical device. *See also id*. 6:29–33 ("Virtualization software in the virtual machine server abstracts the underlying hardware by creating an interface to virtual machines, which represent virtualized resources such as processors, physical memory, network connections, and block devices."). Contrary to IV's proposal, there is no support in the intrinsic evidence for the virtual interface layer to "emulate layers of a networking or storage protocol stack." Accordingly, VMware's proposed construction is entirely consistent with the intrinsic evidence and with how a POSITA would understand the meaning of the disputed terms when reading the '818 patent and should be adopted by the Court.

E. "one or more input/output virtualization modules comprising computer-readable instructions operative to cause the one or more processors to" performs functions terms ('818 patent claim 17)

Plaintiff concedes that VMware's proposal should "succeed" if the Court finds that the term "'one or more input/output virtualization modules' is a nonce term that would not be understood by a PHOSITA as denoting structure." Dkt. No. 53 at 43. VMware has made this showing. Under *Williamson v. Citrix* the term "module" is recognized as a nonce word that does not connote structure. 792 F.3d 1339, 1350 (Fed. Cir. 2015). Furthermore, the prefix "input/output virtualization" adds no known structure; nor does the specification describe a structural component for this term. *Williamson,* 792 F.3d at 1351 ("finding means-plus-function when there was "nothing in the specification or prosecution history that might lead us to construe that expression as the name of a sufficiently definite structure"). Nor does the suffix of "comprising computer-readable instructions operative to cause the one or more processors to" add any structural component. *Glob. Equity Mgmt (SA) Pty. Ltd. v. Expedia, Inc.*, No. 2:16-cv-00095-RWS-RSP, 2016 WL 7416132, at *29 (E.D. Tex. Dec. 22, 2016) ("the 'program code for configuring ...' term is governed by § 112, ¶ 6."); *Personal Audio, LLC v. Apple, Inc*., No. 9:09-cv-111, 2011 WL 11757163, at *21 (E.D. Tex. Jan. 31, 2011) ("'processor' cannot describe sufficient structure").

IV's reliance on *Zeroclick, LLC v. Apple Inc*., as support for its position is misplaced. In *ZeroClick*, the court found that the term "user interface code" was not a means plus function term. In so finding, the court found that the term was a "specific reference[] to conventional graphical user interface programs or code, existing in prior art at the time of the inventions." 891 F.3d 1003, 1008 (Fed. Cir. 2018). The Court bolstered its opinion by observing that the specification described conventional graphical user interface code. *Id.* at 1009. In contrast, in the present case, the term "input/output virtualization modules" is not a reference to a conventional structure existing in the prior art, nor does the specification discuss a conventional "input/output virtualization modules" or even recite the term.

IV's remaining arguments are misplaced. For example, IV argues that for the claimed function of "maintain a connection, over a network fabric, to a virtual storage network interface layer of an application server, wherein the virtual storage network interface layer is associated with a virtual storage node identifier" there is sufficient structure recited in the claim of "a network fabric, a virtual storage network interface layer, and a virtual storage node identifier." Dkt. No. 53 at 42. But, IV fails to realize that none of these components perform the "maintain" function. *See Inventio AG v. ThyssenKrupp Elevator Am. Corp*., 649 F.3d 1350, 1356 (Fed. Cir. 2011) (analyzing whether there is "sufficient structure for performing [the claimed] function"). Instead, the "input/output virtualization modules" are recited in the claim as performing the "maintain" function; the components that IV identified are merely the endpoints being connected.

The same rational supports rejecting each of IV's arguments for these terms. For example, the "physical storage network interface" doesn't perform "enforce" function,[23] it is merely the

[23] *See, e.g.,* Dkt. No. 53-4 at 2–3 ("**enforce** a hierarchical token bucket resource allocation of bandwidth **across the physical storage network interface**") (emphasis added).

medium across which the functions are performed. Nor do the presence of other allegedly structural terms dispersed throughout the clams as objects of the functions, or "additional elements external to the disputed elements" (Dkt. No. 53 at 42), bear any relevance to the structure for performing the "maintain", "present", "enforce", etc. functions recited in claim 17—instead, the claim explicitly recites that it is the input/output virtualization module that causes one or more processors to perform these functions.

Finally, IV's alternative identification of structure disclosed in the specification is unsupported. *See* Dkt. No. 53 at 43. IV has made no attempt to show *any* disclosure of these alleged structures performing the claimed functions. Instead, IV merely proffered a lengthy string-cite without substantive analysis. *Id.* (citing to the '818 patent at 2:9–18; 3:9–11; 3:15–30; 3:43–53; 3:60–4:13; 4:27–5:32, 5:55–6:42, 7:26–67; 8:1–12:28; 12:30–13:4; 13:6–14:29; 14:65–15:23;15:46–17:19; Figs. 1–4; and Figs. 11–13). Dkt. No. 53 at 43. Regardless, VMware's expert has reviewed these sections and found that a PHOSITA would not have seen structure in these sections for performing the claimed functions. Dkt. No. 54-2 (Snoeren Decl.) at 52. As such, this claim is indefinite.

VII. CONCLUSION

For the reasons stated herein, VMware respectfully requests the Court adopt its proposed constructions for the disputed terms and phrases.

Dated: March 27, 2020

Respectfully submitted,

WINSTON & STRAWN LLP

/s/ *Katherine Vidal*
Katherine Vidal
Admitted *Pro Hac Vice*
KVidal@winston.com
Michael R. Rueckheim
Texas State Bar No. 24081129
MRuekheim@winston.com
WINSTON & STRAWN LLP
275 Middlefield Road, Suite 205
Menlo Park, CA 94025
Telephone: (650) 858-6500
Facsimile: (650) 858-6559

Thomas M. Melsheimer
Texas State Bar No. 13922550
M. Brett Johnson
Texas State Bar No. 00790975
MBJohnson@winston.com
Michael A. Bittner
MBittner@winston.com
Admitted *Pro Hac Vice*
Winston & Strawn LLP
2121 N. Pearl St., 9th Floor
Dallas, TX 75201
Telephone: (214) 453-6500
Facsimile: (214) 453-6400

Vivek V. Krishnan
Admitted *Pro Hac Vice*
VKrishnan@winston.com
DaWanna McCray
Admitted *Pro Hac Vice*
DMcCray@winston.com
Winston & Strawn LLP
35 W. Wacker Drive
Chicago, IL 60601
Telephone: (312) 558-5600
Facsimile: (312) 558-5700

William M. Logan
Texas State Bar No. 24106214
WLogan@winston.com
Winston & Strawn LLP
1111 Louisiana Street
Houston, TX 77002
Telephone: (713) 651-2766
Facsimile: (713) 651-2700

ATTORNEYS FOR DEFENDANT VMWARE, INC.

CERTIFICATE OF SERVICE

I hereby certify that, on March 27, 2020, the foregoing document was electronically filed with the Clerk of Court using the Court's CM/ECF system which will send notification of such filing to all counsel of record, including counsel of record for Plaintiffs Intellectual Ventures I LLC and Intellectual Ventures II LLC.

/s/ *Katherine Vidal*
Katherine Vidal